arXiv:1704.07040v1 [math.ST] 24 Apr 2017

Bootstrapping for multivariate linear regression models

Daniel J. Eck
School of Statistics, University of Minnesota, Minneapolis, MN, United States

Abstract

The multivariate linear regression model is an important tool for investigating relationships between several response variables and several predictor variables. The primary interest is in inference about the unknown regression coefficient matrix. We propose multivariate bootstrap techniques as a means for making inferences about the unknown regression coefficient matrix. These bootstrapping techniques are extensions of those developed in Freedman (1981), which are only appropriate for univariate responses. Extensions to the multivariate linear regression model are made without proof. We formalize this extension and prove its validity.

Key Words: Multivariate Bootstrap; Multivariate Linear Regression Model; Residual Bootstrap

1 Introduction

The linear regression model is an important and useful tool in many statistical analyses for studying the relationship among variables. Regression analysis is primarily used for predicting values of the response variable at interesting values of the predictor variables, discovering the predictors that are associated with the response variable, and estimating how changes in the predictor variables affects the response variable (Weisberg, 2005). The standard linear regression methodology assumes that the response variable is a scalar. However, it may be the case that one is interested in investigating multiple response variables simultaneously. One could perform a regression analysis on each response separately in this setting. Such an analysis would fail to detect associations between responses. Regression settings where associations of multiple responses is of interest require a multivariate linear regression model for analysis.

Bootstrapping techniques are well understood for the linear regression model with a univariate response (Bickel and Freedman, 1981; Freedman, 1981). In particular, theoretical justification for the residual bootstrap as a way to estimate the variability of the ordinary least squares (OLS) estimator of the regression coefficient vector in this model has been developed (Freedman, 1981). Theoretical extensions of residual bootstrap techniques appropriate for the multivariate linear regression model have not been formally introduced. The existence of such an extension is stated without proof and rather implicitly in subsequent works (Freedman and Peters, 1984; Diaconis and Efron, 1983). In this article we show that the bootstrap procedures in Freedman (1981) provide consistent estimates of the variability of the OLS estimator of the regression coefficient matrix in the multivariate linear regression model. Our proof technique follows similar logic as Freedman (1981). The generality of the bootstrap theory developed in Bickel and Freedman (1981) provide the tools required for our extension to the multivariate linear regression model.

2 Bootstrap for the multivariate linear regression model

The multivariate linear regression is

$$Y_i = \beta X_i + \varepsilon_i, \quad (i = 1, ..., n), \tag{1}$$

where $Y_i \in \mathbb{R}^r$ and $r > 1$ in order to have an interesting problem, $\beta \in \mathbb{R}^{r \times p}$, $X_i \in \mathbb{R}^p$, and the ε_i's are errors having mean zero and variance Σ where $\Sigma > 0$. It is assumed that separate realizations from the model (1)

are independent. We further define $\mathbb{X} \in \mathbb{R}^{n \times p}$ as the design matrix with rows X_i^T, $\mathbb{Y} \in \mathbb{R}^{n \times r}$ is the matrix of responses with rows Y_i^T, and $\varepsilon \in \mathbb{R}^{n \times r}$ is the matrix of all errors with rows ε_i^T. The OLS estimator of β in model (1) is $\hat{\beta} = \mathbb{Y}^T\mathbb{X}(\mathbb{X}^T\mathbb{X})^{-1}$. We let $\widehat{\varepsilon} \in \mathbb{R}^{n \times r}$ denote the matrix of residuals consisting of rows $\widehat{\varepsilon}_i^T = (Y_i - \hat{\beta}X_i)^T$. The multivariate linear regression model assumed here is slightly different than the traditional multivariate linear regression model. The traditional model makes the additional assumptions that the errors are normally distributed and the design matrix $\mathbb{X}$ is fixed.

We consider two bootstrap procedures that consistently estimate the asymptotic variability of $\text{vec}(\hat{\beta})$ under different assumptions placed upon the model (1). The first bootstrap procedure is appropriate when the design matrix $\mathbb{X}$ is assumed to be fixed and the errors are homoscedastic. In this setup, residuals are resampled. The second bootstrap procedure is appropriate when $(X_i^T, \varepsilon_i^T)^T$ are realizations from a joint distribution. In this setup, cases $(X_i^T, Y_i^T)^T$ are resampled. It is known that bootstrapping under these setups provides a consistent estimator of the variability of $\text{var}(\hat{\beta})$ in model (1) when $r = 1$ (Freedman, 1981). We now provide the needed extensions.

2.1 Fixed design

We first establish the residual bootstrap of Freedman (1981) when $\mathbb{X}$ is assumed to be a fixed design matrix. Resampled, starred, data is generated by the model

$$\mathbb{Y}^* = \mathbb{X}\hat{\beta}^T + \varepsilon^*, \tag{2}$$

where $\varepsilon^* \in \mathbb{R}^{n \times r}$ is the matrix of errors with rows being independent with common distribution $\widehat{F}_n$, the empirical distribution of the residuals, centered at their mean, from the original dataset. Now $\hat{\beta}^* = {\mathbb{Y}^*}^T\mathbb{X}(\mathbb{X}^T\mathbb{X})^{-1}$ is the OLS estimator of β from the starred data. This process is performed a total of B times with a new estimator $\hat{\beta}^*$ computed from (2) at each iteration. We then estimate the variability of $\text{vec}(\hat{\beta})$ with

$$\text{var}^*\left\{\text{vec}(\hat{\beta})\right\} = (B-1)^{-1}\sum_{b=1}^{B}\left\{\text{vec}(\hat{\beta}_b^*) - \text{vec}(\bar{\beta}^*)\right\}\left\{\text{vec}(\hat{\beta}_b^*) - \text{vec}(\bar{\beta}^*)\right\}^T$$

where $\hat{\beta}_b^*$ is the residual bootstrap estimator of β at iteration b and $\bar{\beta}^* = B^{-1}\sum_{b=1}^{B}\hat{\beta}_b^*$.

Before the theoretical justification of the residual bootstrap is formally given, some important quantities are stated. The residuals from the regression (2) are $\varepsilon^* = \mathbb{Y}^* - \mathbb{X}{\hat{\beta}^*}^T$. The variance-covariance matrix Σ in model (1) is then estimated by

$$\widehat{\Sigma} = n^{-1}\sum_{i=1}^{n}\widehat{\varepsilon}_i\widehat{\varepsilon}_i^T - \hat{\mu}^2, \qquad \hat{\mu}^2 = \left(n^{-1}\sum_{i=1}^{n}\widehat{\varepsilon}_i\right)\left(n^{-1}\sum_{i=1}^{n}\widehat{\varepsilon}_i\right)^T.$$

Likewise, the variance-covariance estimate from the starred data is

$$\widehat{\Sigma}^* = n^{-1}\sum_{i=1}^{n}\widehat{\varepsilon}_i^*{\widehat{\varepsilon}_i^*}^T - {\hat{\mu}^*}^2, \qquad {\hat{\mu}^*}^2 = \left(n^{-1}\sum_{i=1}^{n}\widehat{\varepsilon}_i^*\right)\left(n^{-1}\sum_{i=1}^{n}\widehat{\varepsilon}_i^*\right)^T.$$

Theorem 1 provides bootstrap asymptotics for the regression model (1). It extends Theorem 2.2 of Freedman (1981) to the multivariate setting.

Theorem 1. *Assume the regression model* (1) *where the errors have finite fourth moments. Suppose that* $n^{-1}\mathbb{X}^T\mathbb{X} \to \Sigma_X > 0$. *Along almost all sample paths* $Y_1, ..., Y_n$, *as* n *tends to* ∞,

a) the distribution of $\sqrt{n}\left\{vec(\hat{\beta}^*) - vec(\hat{\beta})\right\}$ *converges weakly to a normal distribution with mean 0 and variance-covariance matrix given by* $\Sigma_X^{-1} \otimes \Sigma$.

b) the sequence $\widehat{\Sigma}^*$ *converges to* Σ *in probability.*

c) the distribution of $\left\{(\mathbb{X}^T\mathbb{X})^{1/2} \otimes \widehat{\Sigma}^{*^{-1/2}}\right\}\left\{vec(\hat{\beta}^*) - vec(\hat{\beta})\right\}$ *converges to a standard normal distribution in* $\mathbb{R}^{rp}$.

The proof of Theorem 1, along with the details of several necessary lemmas and theorems, are included in the theoretical details section. Theorem 1 establishes the multivariate analogue for the residual bootstrap. This theorem shows that standard error estimation of the estimated β matrix obtained through bootstrapping, is $\sqrt{n}$-consistent. Now let $f : \mathbb{R}^{rp} \to \mathbb{R}^k$ be a differentiable function. Then the conclusions of Theorem 1 can be applied to establish a multivariate delta method based on estimates obtained via the residual bootstrap. This immediately follows from a first order Taylor expansion and some algebra arriving at

$$\begin{aligned}\sqrt{n}\left[f\left\{\text{vec}(\hat{\beta}^*)\right\} - f\left\{\text{vec}(\hat{\beta})\right\}\right] \\ = \nabla f\left\{\text{vec}(\hat{\beta})\right\}\sqrt{n}\left\{\text{vec}(\hat{\beta}^*) - \text{vec}(\hat{\beta})\right\} + O_p(n^{-1/2})\end{aligned} \tag{3}$$

Therefore (3) converges to a normal distribution with mean zero and variance given by

$$\nabla f\left\{\text{vec}\left(\beta\right)\right\}\left(\Sigma_X^{-1} \otimes \Sigma\right)\nabla^T f\left\{\text{vec}\left(\beta\right)\right\}$$

as $n \to \infty$.

2.2 Random design and heteroskedasticity

In this section we assume that the X_is in model (1) are realizations of a random variable X. The regression coefficient matrix β now takes the form $\beta = E(YX^T)\Sigma_X^{-1}$ where $\Sigma_X = E(XX^T)$ and it is assumed that Σ_X is positive definite. Now that X is stochastic, there may be some association between X and the errors ε. The possibility of heteroskedasticity means that we need to alter the bootstrap procedure outlined in the previous section in order to consistently estimate the variability of vec$(\hat{\beta})$.

It is assumed that the data vectors $(X_i^T, Y_i^T)^T \in \mathbb{R}^{p+r}$ are independent, with a common distribution μ and $E(\|(X_i^T, Y_i^T)^T\|^4) < \infty$ where $\|\cdot\|$ is the Euclidean norm. Unlike the fixed design setting, data pairs $(X_i^T, Y_i^T)^T$ are resampled with replacement to form the starred data $(X_i^{*^T}, Y_i^{*^T})^T$, for $i = 1, ..., n$. Given the original sample, $(X_i^T, Y_i^T)^T$, $i = 1, ..., n$, the resampled vectors are independent, with distribution μ_n. Denote $\mathbb{X}^* \in \mathbb{R}^{n\times p}$ and $\mathbb{Y}^* \in \mathbb{R}^{n\times r}$ as the matrix with rows $X_i^{*^T}$ and $Y_i^{*^T}$ respectively. The starred estimator of β obtained from resampling is then $\hat{\beta}^* = \mathbb{Y}^{*^T}\mathbb{X}^*\left(\mathbb{X}^{*^T}\mathbb{X}^*\right)^{-1}$. For every n there is positive probability, albeit low, that $\mathbb{X}^{*^T}\mathbb{X}^*$ is singular, and the probability of singularity decreases exponentially in n. We assume that displayed equation (1.17) in Chatterjee and Bose (2000) holds in order to circumvent singularity in our bootstrap procedure.

The bootstrap is performed a total of B times with a new estimator $\hat{\beta}^*$ computed at each iteration. We then estimate the variability of vec$(\hat{\beta})$ with

$$\text{var}^*\left\{\text{vec}(\hat{\beta})\right\} = (B-1)^{-1}\sum_{b=1}^{B}\left\{\text{vec}(\hat{\beta}_b^*) - \text{vec}(\bar{\beta}^*)\right\}\left\{\text{vec}(\hat{\beta}_b^*) - \text{vec}(\bar{\beta}^*)\right\}^T$$

where $\hat{\beta}_b^*$ is the bootstrap estimator of β at iteration b and $\bar{\beta}^* = B^{-1}\sum_{b=1}^{B}\hat{\beta}_b^*$.

We now show that the variability of vec$(\hat{\beta})$ is estimated consistently by our multivariate bootstrap procedure which resamples cases. Let M be a non-negative definite matrix with entries $M_{jk} = E\left\{\text{vec}(X_i\varepsilon_i^T)_j\text{vec}(X_i\varepsilon_i^T)_k\right\}$ for $j,k = 1,...,rp$ and define $\Delta = \left(\Sigma_X^{-1}\otimes I_r\right)M\left(\Sigma_X^{-1}\otimes I_r\right)$. where $n^{-1}\mathbb{X}^T\mathbb{X} \to \Sigma_X$ a.e. as $n\to\infty$. Then

$$\begin{aligned}\sqrt{n}\text{vec}\left(\hat{\beta}-\beta\right) &= \text{vec}\left\{n^{-1/2}\varepsilon^T\mathbb{X}(n^{-1}\mathbb{X}^T\mathbb{X})^{-1}\right\} \\ &= \left\{(n^{-1}\mathbb{X}^T\mathbb{X})^{-1}\otimes I_r\right\}\text{vec}\left(n^{-1/2}\varepsilon^T\mathbb{X}\right) \to N(0,\Delta).\end{aligned} \tag{4}$$

The next theorem states that $\sqrt{n}\text{vec}\left(\hat{\beta}^*-\hat{\beta}\right)$ is the same as (4). This is an extension of Theorems 3.1 and 3.2 of Freedman (1981) to the multivariate linear regression setting.

Theorem 2. *Assume that $(X_i^T, Y_i^T)^T \in \mathbb{R}^{p+r}$ are independent, with a common distribution μ, $E(\|(X_i^T,Y_i^T)^T\|^4) < \infty$, and $\Sigma_X = E(XX^T)$ is positive definite. Along almost all sample sequences, given $(X_i^T,Y_i^T)^T$, for $i = 1,...,n$, as $n\to\infty$,*

a) $n^{-1}\left(\mathbb{X}^{*T}\mathbb{X}^*\right)$ converges to Σ_X in conditional probability,

b) the conditional law of $\sqrt{n}\left\{vec(\hat{\beta}^*) - vec(\hat{\beta})\right\}$ converges weakly to a Normal random variable that is mean 0 and has variance-covariance matrix Δ,

c) the sequence $\widehat{\Sigma}^*$ converges to Σ in probability.

The proof of Theorem 2, along with necessary lemmas, are included in the theoretical details section.

3 Examples

In this section we provide two simulated examples which show support for our multivariate bootstrap procedures.

3.1 Example 1: fixed design

This example illustrates Theorem 1. We generated data according to the multivariate linear regression model (1) where $Y_i \in \mathbb{R}^3$, $X_i \in \mathbb{R}^2$, and both β and Σ are prespecified. Our goal is to make inference about vec(β) using the residual bootstrap. Bootstrapped standard errors are compared to the true standard deviations, which are the square root of the diagonal elements of $n^{-1}(\mathbb{X}^T\mathbb{X})^{-1}\otimes\Sigma$. Three data sets were generated at different sample sizes and the performance of the multivariate residual bootstrap is assessed. The bootstrap sample size is taken to be $B = 4n$ in each dataset. The results are displayed in Table 1. We see that the standard errors of vec$(\hat{\beta})$ obtained from the residual bootstrap are close to the true standard errors and that the distance between the two shrinks as n increases.

$n = 100$		$n = 500$		$n = 1000$	
$\text{se}^*\{\text{vec}(\hat{\beta})\}$	$\text{se}_{\text{true}}\{\text{vec}(\hat{\beta})\}$	$\text{se}^*\{\text{vec}(\hat{\beta})\}$	$\text{se}_{\text{true}}\{\text{vec}(\hat{\beta})\}$	$\text{se}^*\{\text{vec}(\hat{\beta})\}$	$\text{se}_{\text{true}}\{\text{vec}(\hat{\beta})\}$
0.249	0.281	0.144	0.140	0.091	0.097
0.317	0.336	0.146	0.141	0.093	0.100
0.237	0.251	0.129	0.125	0.080	0.087
0.270	0.301	0.133	0.126	0.082	0.089
0.140	0.154	0.077	0.077	0.050	0.053
0.168	0.184	0.079	0.077	0.054	0.055

Table 1: Comparison of bootstrapped standard errors to the true standard errors across three sample sizes in the fixed design case. The bootstrap sample size is set at $B = 4n$ for each dataset.

3.2 Example 2: random design and heteroskedasticity

This example aims to show support for Theorem 2. We generated data according to the multivariate linear regression model (1) where $Y_i \in \mathbb{R}^3$, $X_i \in \mathbb{R}^2$, and both β and Σ are prespecified. The predictors and errors are generated according to

$$\begin{pmatrix} X_i \\ \varepsilon_i \end{pmatrix} \sim N\left\{ \begin{pmatrix} 0 \\ 0 \end{pmatrix}, \begin{pmatrix} \Sigma_X & \Sigma_{X\varepsilon} \\ \Sigma_{\varepsilon X} & \Sigma \end{pmatrix} \right\},$$

for $i = 1, ..., n$. Our goal is to make inference about $\text{vec}(\beta)$ using the multivariate bootstrap procedure in the random design case. Bootstrapped standard errors are compared to the true standard deviations, which are the square root of the diagonal elements of $(n-1)^{-1}\Sigma_X^{-1} \otimes \left(\Sigma - \Sigma_{\varepsilon X}\Sigma_X^{-1}\Sigma_{X\varepsilon}\right)$. Three data sets were generated at different sample sizes and the performance of the multivariate bootstrap is assessed. The bootstrap sample size is taken to be $B = 4n$ in each dataset. The results are displayed in Table 2. We see that the standard errors of $\text{vec}(\hat{\beta})$ obtained from the bootstrap are close to the true standard errors.

$n = 100$		$n = 500$		$n = 1000$	
$\text{se}^*\{\text{vec}(\hat{\beta})\}$	$\text{se}_{\text{true}}\{\text{vec}(\hat{\beta})\}$	$\text{se}^*\{\text{vec}(\hat{\beta})\}$	$\text{se}_{\text{true}}\{\text{vec}(\hat{\beta})\}$	$\text{se}^*\{\text{vec}(\hat{\beta})\}$	$\text{se}_{\text{true}}\{\text{vec}(\hat{\beta})\}$
0.340	0.310	0.131	0.138	0.101	0.098
0.257	0.275	0.106	0.123	0.082	0.087
0.146	0.159	0.068	0.071	0.046	0.050
0.333	0.310	0.138	0.138	0.104	0.098
0.225	0.275	0.124	0.123	0.086	0.087
0.135	0.159	0.071	0.071	0.051	0.050

Table 2: Comparison of bootstrapped standard errors to the true standard errors across three sample sizes in the random design case with heteroskedasticity. The bootstrap sample size is set at $B = 4n$ for each dataset.

4 Theoretical details

Before we present our proof of Theorems 1 and 2, we motivate the Mallows metric as a central tool for our proof technique. The Mallows metric for probabilities in $\mathbb{R}^p$, relative to the Euclidean norm was the driving

force needed to establish the validity of the residual bootstrap approximation in the context of univariate regression (Bickel and Freedman, 1981; Freedman, 1981). The Mallows metric, relative to the Euclidean norm, for two probability measures μ, ν in $\mathbb{R}^p$, denoted $d_l^p(\mu, \nu)$, is the infimum of $E^{1/l}\left(\|U - V\|^l\right)$ over all joint distributions of random vectors U and V, where U has distribution μ and V has distribution ν. Properties of the Mallows metric are developed for random variables on separable Banach spaces of finite dimension (Bickel and Freedman, 1981). Since $\mathbb{R}^k$ is indeed a separable Banach space for a natural number k, the theory in Bickel and Freedman (1981) applies to our case. In the present paper, we use the Mallows metric when $r > 1$ to prove that the residual bootstrap can be used to estimate the variability of $\text{vec}(\hat{\beta})$ consistently.

4.1 Fixed design

Let $\Psi_n(F)$ be the distribution function of $\sqrt{n}\left\{\text{vec}(\hat{\beta}) - \text{vec}(\beta)\right\}$ where F is the law of the errors ε so that $\Psi_n(F)$ is a probability measure on $\mathbb{R}^{rp}$. Let G be an alternate law of the errors, where it is assumed that G is mean-zero with finite variance $\Sigma_G > 0$. In applications, G will be the centered empirical distribution of the residuals.

Theorem 3. $\left[d_2^{rp}\left\{\Psi_n(F), \Psi_n(G)\right\}\right]^2 \leq nr\,\text{tr}\left\{(\mathbb{X}^T\mathbb{X})^{-1}\right\}\left\{d_2^r(F,G)\right\}^2$.

Proof. Let $B = \mathbb{X}(\mathbb{X}^T\mathbb{X})^{-1}$. Then $\Psi_n(F)$ is the law of $\sqrt{n}\varepsilon_n^T(F)B$ where $\varepsilon_n(F)$ is the matrix with n rows of independent random variables ε, having common law F. $\Psi_n(G)$ can be thought of similarly. Observe that $B^TB = (\mathbb{X}^T\mathbb{X})^{-1}$. Then, from Lemma 8.9 in Bickel and Freedman (1981), we see that

$$\begin{aligned}
\left[d_2^{rp}\left\{\Psi_n(F), \Psi_n(G)\right\}\right]^2 &= \left(d_2^{rp}\left[\text{vec}\{\varepsilon_n^T(F)B\}, \text{vec}\{\varepsilon_n^T(G)B\}\right]\right)^2 \\
&= \left(d_2^{rp}\left[(B^T \otimes I_r)\text{vec}\{\varepsilon_n^T(F)\}, (B^T \otimes I_r)\text{vec}\{\varepsilon_n^T(G)\}\right]\right)^2 \\
&\leq n\,\text{tr}\left\{(B^T \otimes I_r)(B^T \otimes I_r)^T\right\}\left\{d_2^r(F,G)\right\}^2 = n\,\text{tr}\left\{(B^T \otimes I_r)(B \otimes I_r)\right\}\left\{d_2^r(F,G)\right\}^2 \\
&= n\,\text{tr}\left(B^TB \otimes I_r\right)\left\{d_2^r(F,G)\right\}^2 = n\,\text{tr}\left\{(\mathbb{X}^T\mathbb{X})^{-1} \otimes I_r\right\}\left\{d_2^r(F,G)\right\}^2 \\
&= nr\,\text{tr}\left\{(\mathbb{X}^T\mathbb{X})^{-1}\right\}\left\{d_2^r(F,G)\right\}^2,
\end{aligned}$$

which is our desired conclusion. □

With Theorem 3 we can bound the distance between the sample dependent distribution functions $\Psi_n(F)$ and $\Psi_n(G)$ by the distance between their underlying laws. As in Freedman (1981), we proceed with F_n as the empirical distribution function of $\varepsilon_1, ..., \varepsilon_n$. Let $\widetilde{F}_n$ be the empirical distribution of the residuals $\widehat{\varepsilon}_1, ..., \widehat{\varepsilon}_n$ from the original regression, and let $\widehat{F}_n$ be $\widetilde{F}_n$ centered at its mean $\hat{\mu} = n^{-1}\sum_{i=1}^n \widehat{\varepsilon}_i$. Since $\widehat{\varepsilon} = \mathbb{Y} - \mathbb{X}\hat{\beta}^T$, we have $\widehat{\varepsilon} - \varepsilon = -\mathcal{P}\varepsilon$ where $\mathcal{P}$ is the projection into the column space of $\mathbb{X}$.

Lemma 1. $E^2\left\{d_2^r(\widetilde{F}_n, F_n)\right\} \leq p\,\text{tr}(\Sigma)/n$.

Proof. From the definition of the Mallows metric we have

$$\begin{aligned}
\left\{d_2^r(\widetilde{F}_n, F_n)\right\}^2 &\leq n^{-1}\sum_{i=1}^n \|\widehat{\varepsilon}_i - \varepsilon_i\|^2 = n^{-1}\,\text{tr}\left\{(\widehat{\varepsilon} - \varepsilon)^T(\widehat{\varepsilon} - \varepsilon)\right\} \\
&= n^{-1}\,\text{tr}\left(\varepsilon^T\mathcal{P}\varepsilon\right).
\end{aligned}$$

From linearity of the expectation with respect to the trace operator,

$$E\left\{\operatorname{tr}\left(\varepsilon^T\mathcal{P}\varepsilon\right)\right\} = \operatorname{tr}\left\{E\left(\varepsilon^T\mathcal{P}\varepsilon\right)\right\} = \operatorname{tr}\left\{\mathcal{P}E\left(\varepsilon\varepsilon^T\right)\right\} \leq \operatorname{tr}\left(\mathcal{P}\right)\operatorname{tr}\left(\Sigma\right) = p\operatorname{tr}\left(\Sigma\right)$$

and this completes the proof. □

Lemma 2. $E^2\left\{d_2^r(\widehat{F}_n, F_n)\right\} \leq (p+1)\operatorname{tr}(\Sigma)/n.$

Proof. From Lemma 8.8 in Bickel and Freedman (1981) we have

$$\begin{aligned} d_2^r(\widehat{F}_n, F_n)^2 &= d_2^r\{\widetilde{F}_n - E(\widetilde{F}_n), F_n - E(F_n)\}^2 + \|E(F_n)\|^2 \\ &= d_2^r(\widetilde{F}_n, F_n)^2 - \|E(\widetilde{F}_n) - E(F_n)\|^2 + \|E(F_n)\|^2 \\ &\leq d_2^r(\widetilde{F}_n, F_n)^2 + \|n^{-1}\sum_{i=1}^n \varepsilon_i\|^2 \end{aligned}$$

with the empirical distribution functions F_n,$\widetilde{F}_n$, and $\widehat{F}_n$ used as random variables in the application of Lemma 8.8 in Bickel and Freedman (1981). We see that

$$E\left(\|n^{-1}\sum_{i=1}^n \varepsilon_i\|^2\right) = n^{-2}\left\{E\left(\sum_{i=1}^n \varepsilon_i^T\varepsilon_i + \sum_{i\neq j}\varepsilon_i^T\varepsilon_j\right)\right\} = n^{-1}\left\{E(\varepsilon_1^T\varepsilon_1)\right\} = n^{-1}\operatorname{tr}\left(\Sigma\right).$$

Our conclusion follows from Lemma 1. □

These results imply the validity of the bootstrap approximation, in probability, for the model (1) if we assume that $n^{-1}\mathbb{X}^T\mathbb{X} \to \Sigma_X > 0$. From Theorem 3,

$$E\left[d_2^{rp}\{\Psi_n(\widehat{F}_n), \Psi_n(F)\}\right] \leq nr\operatorname{tr}\left\{(\mathbb{X}^T\mathbb{X})^{-1}\right\} d_2^r(\widehat{F}_n, F)$$

and because of the metric properties of $d_2^r(\cdot,\cdot)$

$$\frac{1}{2}d_2^r(\widehat{F}_n, F)^2 \leq d_2^r(\widehat{F}_n, F_n)^2 + d_2^r(F_n, F)^2$$

where Lemma 2 shows that the first term on the right hand side converges in probability to 0; the second term on the right hand side converges to 0 in probability by Lemma 8.4 of (Bickel and Freedman, 1981) where the separable Banach space is taken to be $\mathbb{R}^r$. The next results are special cases of Lai et. al. (1979) which are adapted from Freedman (1981) to the multivariate setting. We let ε_j, $j = 1, ..., r$, be the column of ε corresponding to the errors of response Y_j.

Lemma 3. $n^{-1}\mathbb{X}^T\varepsilon \to 0$ *almost everywhere and* $\hat{\beta} \to \beta$ *almost everywhere.*

Proof. Let B_j be the jth column of ε. Then $n^{-1}\mathbb{X}^T\varepsilon \in \mathbb{R}^{p\times r}$ with columns $n^{-1}\mathbb{X}^T\varepsilon$. Lemma 2.3 of Freedman (1981) state that $n^{-1}\mathbb{X}^TB_j \to 0$ almost everywhere for any particular $j = 1, ..., r$. Therefore $n^{-1}\mathbb{X}^T\varepsilon \to 0$ almost everywhere. A similar argument verifies our second result. □

Lemma 4. $n^{-1}\operatorname{tr}\left\{(\widehat{\varepsilon}-\varepsilon)^T(\widehat{\varepsilon}-\varepsilon)\right\} \to 0$ *almost everywhere.*

Proof. A similar argument to that of Lemma 2.4 in Freedman (1981) gives

$$\begin{aligned} n^{-1}\operatorname{tr}\left\{(\widehat{\varepsilon}-\varepsilon)^T(\widehat{\varepsilon}-\varepsilon)\right\} &= n^{-1}\operatorname{tr}\left\{\varepsilon^T\mathbb{X}(\mathbb{X}^T\mathbb{X})^{-1}\mathbb{X}^T\varepsilon\right\} \\ &= \operatorname{tr}\left\{\left(n^{-1}\varepsilon^T\mathbb{X}\right)\left(n^{-1}\mathbb{X}^T\mathbb{X}\right)^{-1}\left(n^{-1}\mathbb{X}^T\varepsilon\right)\right\}. \end{aligned}$$

The center term converges to $\Sigma_X > 0$ and the left and right terms converge to 0 almost everywhere by Lemma 3. Our result follows. □

Lemma 5. *$d_2^r(\widehat{F}_n, F_n) \to 0$ almost everywhere and $d_2^r(\widehat{F}_n, F) \to 0$ almost everywhere.*

Proof. From the arguments in the proofs of Lemmas 1 and 2 we have that

$$\begin{aligned} d_2^r(\widehat{F}_n, F_n) &= d_2^r(\widetilde{F}_n, F_n)^2 - \|E(\widetilde{F}_n) - E(F_n)\|^2 + \|E(F_n)\|^2 \\ &= \|n^{-1}\sum_{i=1}^n \varepsilon_i\|^2 - \|n^{-1}\sum_{i=1}^n (\widehat{\varepsilon}_i - \varepsilon_i)\|^2 + d_2^r(\widetilde{F}_n, F_n) \\ &\leq \|n^{-1}\sum_{i=1}^n \varepsilon_i\|^2 + n^{-1}\operatorname{tr}\left\{(\widehat{\varepsilon}-\varepsilon)^T(\widehat{\varepsilon}-\varepsilon)\right\} \end{aligned}$$

which converges to 0 almost everywhere by Lemma 4. Therefore the first convergence result holds. From the metric properties of the Mallows metric we have that

$$\frac{1}{2}d_2^r(\widehat{F}_n, F)^2 \leq d_2^r(\widehat{F}_n, F_n)^2 + d_2^r(F_n, F)^2.$$

Our second convergence result follows from the first convergence result and Lemma 8.4 of Bickel and Freedman (1981). □

Lemma 6. *Let u_i and v_i, $i = 1, ..., n$, be $r \times 1$ vectors. Let*

$$\bar{u} = n^{-1}\sum_{i=1}^n u_i, \quad \textit{and} \quad s_u^2 = n^{-1}\sum_{i=1}^n (u_i - \bar{u})(u_i - \bar{u})^T$$

and similarly for v. Then

$$\|s_u^2 - s_v^2\|_F^2 \leq \|n^{-1}\sum_{i=1}^n (u_i - v_i)(u_i - v_i)^T\|_F^2$$

where $\|\cdot\|_F$ is the Frobenius norm.

Proof. We have

$$\begin{aligned} \|s_u^2 - s_v^2\|_F^2 &= \sum_{j=1}^n\sum_{k=1}^n |n^{-1}\sum_{i=1}^n (u_i - \bar{u})_j(u_i - \bar{u})_k^T - n^{-1}\sum_{i=1}^n (v_i - \bar{v})_j(v_i - \bar{v})_k^T|^2 \\ &\leq \sum_{j=1}^n\sum_{k=1}^n |n^{-1}\sum_{i=1}^n (u_i - v_i)_j(u_i - v_i)_k|^2 \\ &= \|n^{-1}\sum_{i=1}^n (u_i - v_i)(u_i - v_i)^T\|_F^2, \end{aligned}$$

where the inequality follows from (Freedman, 1981, Lemma 2.7). □

The proof of Theorem 1 is now given.

Proof. Exchange $\widehat{F}_n$ for G in Theorem 3 and observe that

$$d_2^{rp}\left\{\Psi_n(F), \Psi_n(\widehat{F}_n)\right\} \leq nr \operatorname{tr}\left\{(\mathbb{X}^T\mathbb{X})^{-1}\right\} d_2^r(F, \widehat{F}_n)^2.$$

From Lemma 5 we know that $d_2^r(F, \widehat{F}_n)^2 \to 0$ almost everywhere. Our result for part a) follows since F is mean-zero normal with variance $\Sigma_X^{-1} \otimes \Sigma$. We now show that part b) holds. First, we need to establish that $\widehat{\Sigma} \to \Sigma$ almost everywhere. To see this, introduce

$$\Sigma_n = n^{-1}\sum_{i=1}^n \varepsilon_i\varepsilon_i^T - \left(n^{-1}\sum_{i=1}^n \varepsilon_i\right)\left(n^{-1}\sum_{i=1}^n \varepsilon_i\right)^T.$$

Clearly, $\Sigma_n \to \Sigma$ almost everywhere. Let $C_n = n^{-1}\sum_{i=1}^n (\hat{\varepsilon}_i - \varepsilon_i)(\hat{\varepsilon}_i - \varepsilon_i)^T$. We have,

$$\begin{aligned}
\|\widehat{\Sigma} - \Sigma_n\|_F^2 \leq \|C_n\|_F^2 &= \operatorname{tr}(C_nC_n) \leq \operatorname{tr}^2(C_n) \\
&= \operatorname{tr}^2\left\{n^{-1}\sum_{i=1}^n (\widehat{\varepsilon}_i - \varepsilon_i)^T(\widehat{\varepsilon}_i - \varepsilon_i)\right\} \\
&= \operatorname{tr}^2\left\{n^{-1}(\widehat{\varepsilon} - \varepsilon)^T(\widehat{\varepsilon} - \varepsilon)\right\} \to 0
\end{aligned}$$

almost everywhere where the first inequality follows from Lemma 6 with $\widehat{\Sigma}_n$ and Σ_n taking the place of s_u^2 and s_v^2 respectively, the second inequality follows from the fact that C_n is positive definite almost everywhere, and the convergence follows from Lemma 4.

Let $D_n = E\left(\|\widehat{\Sigma}_n^* - \Sigma_n^*\|_F \mid Y_1, ..., Y_n\right)$. From Lemma 6 and the proof of Lemma 1 we see that,

$$\begin{aligned}
D_n &\leq E\left\{\|n^{-1}\sum_{i=1}^n \left(\widehat{\varepsilon}_i^* - \varepsilon_i^*\right)\left(\widehat{\varepsilon}_i^* - \varepsilon_i^*\right)^T\|_F \mid Y_1, ..., Y_n\right\} \\
&\leq E\left[\operatorname{tr}\left\{n^{-1}(\widehat{\varepsilon}^* - \varepsilon^*)^T(\widehat{\varepsilon}^* - \varepsilon^*)\right\} \mid Y_1, ..., Y_n\right] \\
&\leq p \operatorname{tr}\left(\widehat{\Sigma}\right)/n
\end{aligned}$$

where the last inequality follows from the argument that proves Lemma 1 applied to the starred data, and $p \operatorname{tr}\left(\widehat{\Sigma}\right)/n \to 0$ almost everywhere. It remains to show that $\widehat{\Sigma}_n^*$ converges to Σ. Conditional on $Y_1, ..., Y_n$,

$$\begin{aligned}
&d_2^{r(r+1)/2}\left\{n^{-1}\sum_{i=1}^n \operatorname{vech}(\varepsilon_i^*\varepsilon_i^{*T}), n^{-1}\sum_{i=1}^n \operatorname{vech}(\varepsilon_i\varepsilon_i^T)\right\} \\
&\quad \leq d_2^{r(r+1)/2}\left\{\operatorname{vech}(\varepsilon_1^*\varepsilon_1^{*T}), \operatorname{vech}(\varepsilon_1\varepsilon_1^T)\right\}
\end{aligned} \tag{5}$$

by Lemma 8.6 in Bickel and Freedman (1981). Now ε^* has conditional distribution $\widehat{F}_n$ and ε has law F and Lemma 5 gives $d_2^r\left(\widehat{F}_n, F\right) \to 0$ almost everywhere. We now show that $d_1\left\{\operatorname{vech}(\varepsilon_1^*\varepsilon_1^{*T}), \operatorname{vech}(\varepsilon_1\varepsilon_1^T)\right\} \to 0$ almost everywhere by Lemma 8.5 of Bickel and Freedman (1981) with $\phi(x) = \operatorname{vech}\left(xx^T\right)$ where $x \in \mathbb{R}^r$. To do this, we show that K can be chosen so that $\|\phi(x)\|_1 \leq K(1 + \|x\|_2^2)$ where $\|\cdot\|_1$ and $\|\cdot\|_2$ are the

$\mathcal{L}^1$ and $\mathcal{L}^2$ norms respectively. From the definition of the Euclidean norm, we have $\|x\|_2^2 = \sum_{i=1}^r x_i^2$. It is clear that $x_i^2 + x_j^2 \geq 2|x_i x_j|$ for all $i, j = 1, ..., r$. Now, pick $K = \binom{r}{2} + 1$. We see that

$$\begin{aligned} K(1 + \|x\|_2^2) &\geq \Big|\sum_{i=1}^r x_i^2 + \binom{r}{2}\sum_{i=1}^r x_i^2\Big| \geq \Big|\sum_{i=1}^r x_i^2 + \sum_{i\neq j}^r |x_i x_j|\Big| \\ &\geq \Big|\sum_{i\geq j}^r |x_i x_j|\Big| \geq \|\text{vech}\left(xx^T\right)\|_1 = \|\phi(x)\|_1 \end{aligned}$$

A similar argument shows that $1/n \sum_{i=1}^n \varepsilon_i^*$ converges to 0. Part c) follows from both a) and b). □

4.2 Random design and heteroskedasticity

In this section we provide the proof of Theorem 2. Several quantities and lemmas are introduced in order to prove Theorem 2. The logic follows that of (Freedman, 1981, Section 3). Define,

$$\begin{aligned} \Sigma(\mu) &= \int xx^T \mu(dx), \\ \beta(\mu) &= \int yx^T \mu(dx, dy)\Sigma(\mu)^{-1}, \\ \varepsilon(\mu, x, y) &= y - \beta(\mu)x^T. \end{aligned}$$

The next two lemmas are needed to prove Theorem 2.

Lemma 7. *If $d_4^{p+r}(\mu_n, \mu) \to 0$ as $n \to \infty$, then*

a) $\Sigma(\mu_n) \to \Sigma(\mu)$ and $\beta(\mu_n) \to \beta(\mu)$,

b) the μ_n-law of vec$\{\varepsilon(\mu_n, x, y)x^T\}$ converges to the μ-law of vec$\{\varepsilon(\mu, x, y)x^T\}$ in d_2^{rp},

c) the μ_n-law of $\|\varepsilon(\mu_n, x, y)\|^2$ converges to the μ-law of $\|\varepsilon(\mu, x, y)\|^2$ in d_1.

Proof. Part a) immediately follows from (Bickel and Freedman, 1981, Lemma 8.3c).

We use (Bickel and Freedman, 1981, Lemma 8.3a) to verify part b). The weak convergence step is trivial. Now,

$$\begin{aligned} \|\text{vec}\{\varepsilon(\mu_n, x, y)x^T\}\|^2 &= \|\text{vec}\{yx^T - \beta(\mu_n)xx^T\}\|^2 \\ &= \|\text{vec}(yx^T)\|^2 + \|\text{vec}(\beta(\mu_n)xx^T)\|^2 - 2\text{vec}(yx^T)^T\text{vec}\{\beta(\mu_n)xx^T\}. \end{aligned}$$

Let $z = (x^T, y^T)^T$. Part b) follows from, integration with respect to μ_n, part a), and (Bickel and Freedman, 1981, Lemma 8.5) with $\phi(z) = \text{vech}(zz^T)$. The steps involving (Bickel and Freedman, 1981, Lemma 8.5) are similar to those in the proof of Theorem 1.

Part c) follows from the same argument used to prove part b). □

Lemma 8. *$d_4^{p+r}(\mu_n, \mu) \to 0$ a.e. as $n \to \infty$.*

Proof. The steps are the same as those in (Freedman, 1981, Lemma 3.2). □

The proof of Theorem 2 is now given.

Proof. We can write

$$\begin{aligned}
\text{vec}\left\{\sqrt{n}\left(\hat{\beta}^* - \hat{\beta}\right)\right\} &= \text{vec}\left[\sqrt{n}\left\{\mathbb{Y}^{*^T}\mathbb{X}^*(\mathbb{X}^{*^T}\mathbb{X}^*)^{-1} - \hat{\beta}\right\}\right] \\
&= \text{vec}\left[\sqrt{n}\left\{(\varepsilon^* + \mathbb{X}^*\hat{\beta}^T)^T\mathbb{X}^*(\mathbb{X}^{*^T}\mathbb{X}^*)^{-1} - \hat{\beta}\right\}\right] \\
&= \text{vec}\left\{n^{-1/2}\varepsilon^{*^T}\mathbb{X}^*(n^{-1}\mathbb{X}^{*^T}\mathbb{X}^*)^{-1}\right\} \\
&= \text{vec}\left(Z^*W^{*^{-1}}\right) = \left(W^{*^{-1}} \otimes I_r\right)\text{vec}(Z^*)
\end{aligned}$$

where $Z^* = n^{-1/2}\varepsilon^{*^T}\mathbb{X}^*$ and $W^* = n^{-1}\mathbb{X}^{*^T}\mathbb{X}^*$. (Freedman, 1981, Theorem 3.1) shows that the conditional law of W^* converges to Σ_X in probability. This verifies part a).

We now verify part b). From (Bickel and Freedman, 1981, Lemma 8.7), we have

$$d_2^{rp}\left\{\text{vec}(Z^*), \text{vec}(Z)\right\}^2 \leq d_2^{rp}\left\{\text{vec}(X_i^*\varepsilon_i^{*^T}), \text{vec}(X_i\varepsilon_i^T)\right\}^2$$

where the right side goes to 0 a.e. as $n \to \infty$. Lemma 8 states that $\mu_n \to \mu$ a.e. in d_4^{r+p} as $n \to \infty$ and part b) of Lemma 7 implies that the distribution of $\text{vec}(Z^*)$, conditional on (X_i, Y_i), $i = 1, ..., n$, converges to $\text{vec}(Z)$. The random variable $\text{vec}(Z)$ is normally distributed with mean 0 and variance matrix M. Combining this with part a) verifies that the conditional distribution of $\left(W^{*^{-1}} \otimes I_r\right)\text{vec}(Z^*)$ converges to $\left(\Sigma_X^{-1} \otimes I_r\right)\text{vec}(Z)$ as $n \to \infty$. This completes the proof of part b).

Part c) follows from the same argument in the proof of Theorem 1 where Lemmas 8 and 7c combine to show that (5) converges to 0 as $n \to \infty$. Note that $\varepsilon_1^* = Y_1^* - \hat{\beta}X_1^*$ in this argument. This completes the proof. □

5 Acknowledgments

The author would like to thank Karl Oskar Ekvall, Forrest Crawford, Snigdhansu Chatterjee, and Dennis Cook for providing valuable feedback which led to the strengthening of this article.

References

Bickel, P. J. and Freedman, D. A. (1981). Some Asymptotic Theory for the Bootstrap. *The Annals of Statistics*, **9**, 6, 1196-1217.

Chatterjee, S. and Bose, A. (2000). Variance Estimation in High Dimensional Models. *Statistica Sinica*, **10**, 497–515.

Diaconis, P. and Efron, B. (1983). Computer Intensive Methods in Statistics *Scientific America*, **248**.

Freedman, D. A. (1981). Bootstrapping Regression Models. *The Annals of Statistics*, **9**, 6, 1218-1228.

Freedman, D. A. and Peters, S. C. (1984). Bootstrapping a Regression Equation: Some Empirical Results *Journal of the American Statistical Association*, **79**, 97–106.

Lai, T., Robbins, H., and Wei, V. (1979). Strong Consistency of least squares estimated in multiple regression. *Journal of Multivariate Analysis*, **9**, 343-361.

Weisberg, S. (2005). *Applied Linear Regression third edition*. Wiley, New Jersey.